

Detection of Targets in Bandlimited and Spatially Correlated Clutter

Peter Vouras
Radar Division
Naval Research Laboratory
Washington D.C., USA
peter.vouras@nrl.navy.mil

Abstract—This paper describes the preliminary analysis of an unconventional approach for detecting targets in bandlimited clutter that exhibits some degree of spatial correlation from range bin to range bin. The receiver processing chain utilizes a lattice predictor (LP) filter to remove clutter at the output of a Doppler filter bank. Subsequently, an Order-Statistic (OS) Constant False Alarm Rate (CFAR) processing stage is used to threshold the data and a binary integrator creates a binary image of threshold crossings accumulated over several Coherent Processing Intervals (CPIs). A novel technique is applied for removing noise from the binary image using morphological filters, thereby exposing tracks created by any moving targets. Once a track has been detected, the corresponding target is declared present. Test results based on measured clutter data will be presented that suggest the proposed processing techniques may improve detection performance in a clutter environment.

I. INTRODUCTION

Traditional radar techniques for detecting targets in clutter include Pulse Doppler (PD) processing, Space-Time Adaptive Processing (STAP), likelihood ratio tests, and different types of CFAR. Each technique however has drawbacks. PD processing throws away targets whose velocity falls within the clutter notch. STAP requires training data and extensive computational throughput. Likelihood ratio tests and CFAR circuits require some knowledge of the probability distribution of the clutter for optimum performance. Apart from the problem of detecting targets in clutter, there is also the equally challenging problem of mitigating the false alarms due to clutter. For example, by excessively lowering the detection threshold, a target may be detected in clutter, but the resulting false alarms may overwhelm the receiver.

An analysis of measured clutter data recorded from the AN/SPS-49 radar at the Naval Research Laboratory (NRL) Chesapeake Bay Detachment (CBD) reveals that the clutter envelope is at least partially correlated across range bins. In the absence of thermal noise, and in the presence of spatially

correlated clutter, a LP filter should be able to predict the clutter magnitude in an upcoming range bin based on clutter amplitudes in previous range bins [1]. Assuming the prediction is perfect, then any target would appear at the output of the filter as a large error and could be detected. This principle forms the basis for the adaptive processor proposed in this paper.

In the presence of thermal noise, however, it is not sufficient to simply remove the clutter. At the output of the LP filter, the receiver must also threshold the data to minimize false alarms due to noise. Therefore, an OS-CFAR processor is utilized after the LP filter to threshold the data. The threshold crossings from a batch of CPIs are accumulated, passed through a binary integrator, and compared to a second threshold to produce a noisy binary image. In a novel application to radar processing, morphological image filters are used to remove the noise from this image and expose the target tracks. As in track-before-detect (TBD) processing, the presence of these tracks forms the basis for declaring targets present.

The remainder of this paper is organized as follows. In Section II the characteristics of the measured clutter data are described. In Section III the adaptive processor is described in detail and in Section IV sample results are presented based on the measured clutter data. These preliminary results suggest that the two processing techniques based on linear prediction and morphological filters described in this paper may lead to definitive improvements in clutter detection performance and merit further investigation.

II. CHARACTERISTICS OF MEASURED CLUTTER DATA

A. Radar Parameters and Geographic Location

An L-band AN/SPS-49 radar operated by NRL at the CBD was used to record raw in-phase (I) and quadrature (Q) samples at the output of the pulse compression operation, which is sampled by the receiver analog-to-digital (A/D) converter. This data was then analyzed off-line to determine the nature of the clutter received. Fig. 1 shows an aerial view of the radar installation. The radar is located on a cliff overlooking the Chesapeake Bay. The radar antenna sits approximately 160 feet above sea level. Data was recorded over 360° azimuth with the elevation of the main beam peak fixed at 5° above the horizon. The range of the recorded data extends from 10 out to 70 miles from the antenna, which is within the unambiguous range of the radar. Each recorded data file encompasses an 8° azimuth sector and a 12 mile range interval, which corresponds to 97 range bins at the

A/D sampling rate. The duration of each data recording was 22 scans of the antenna which is sufficient to record 3 CPIs of 16 pulses for each scan. The clutter data analyzed for this paper was accumulated from the center beam position of each azimuth sector. All of the recorded data was screened to ensure that the A/D had not saturated during the dwell. The total amount of data analyzed was 14850 CPIs, which corresponds to the returns from 237,600 pulses. The data was recorded over the course of several days, on some of which the weather happened to be rainy or foggy. In short, the recorded data includes combinations of sea clutter, land clutter, and weather clutter.



Figure 1. NRL CBD Installation

Table 1 lists some of the radar system parameters relevant to this study.

TABLE I. RELEVANT RADAR PARAMETERS

Transmitted Frequency	869.77 MHz
Polarization	Horizontal
Pulse Repetition Interval	981 usec
Pulse Repetition Frequency (PRF)	1019.4 Hz
A/D Sampling Rate	666 kHz
Antenna Scan Rate	12 rpm
Width of Range Bin	0.125 mi
Beamwidth (3 dB)	Vertical : 10° Horizontal : 3.3°

B. Characteristics of Measured Clutter

To design a detector for targets in the presence of clutter some characterization of the clutter is necessary. If the probability density of clutter plus signal plus noise is reasonably well separated from the probability density of signal plus noise alone, then a detector can be designed that decides if a target is present based on the outcome of a likelihood ratio test. Otherwise, an approach that separates the target from clutter in the frequency domain using Doppler processing is an option. Fig. 2 illustrates the complementary cumulative distribution of clutter power from all the measured data. The plot shows that almost 20 percent of the recorded range bins contain clutter power

greater than 20 dB. This result suggests that even if the density of clutter amplitudes could be accurately characterized by a functional description, a detection approach based on likelihood ratio tests would most likely fail to detect weak targets because of the large fraction of range cells occupied by large magnitude clutter.

Rather than attempting to detect targets in the time domain, a better detection strategy may be to separate targets at the output of a Doppler filter bank. In this scheme, each target will compete for detection only with the clutter energy within the same Doppler filter. To investigate the utility of this approach, a Doppler filter bank was created by assembling the received radar pulses into CPIs of length 16 and calculating a weighted Discrete Fourier Transform (DFT) at each range cell. Fig. 3 illustrates the complementary cumulative distribution of clutter power at the output of the first filter in the Doppler filter bank. As can be seen, strong clutter amplitudes are still present in this Doppler filter, which will hamper the detection of targets. In general, strong clutter magnitudes were present in all the lower Doppler filters and the problem of separating targets from the clutter remained as severe in the frequency domain as in the time domain. However, the higher Doppler filters, from 3 through 14, contained much lower clutter amplitudes, as seen in Fig. 4. At the output of these Doppler filters, target detection should be much easier and can be accomplished using conventional techniques.




Figure 2. Complementary Cumulative Distribution of Time Domain Clutter Power




Figure 3. Complementary Cumulative Distribution of Clutter Power at Output of Doppler Filter 1

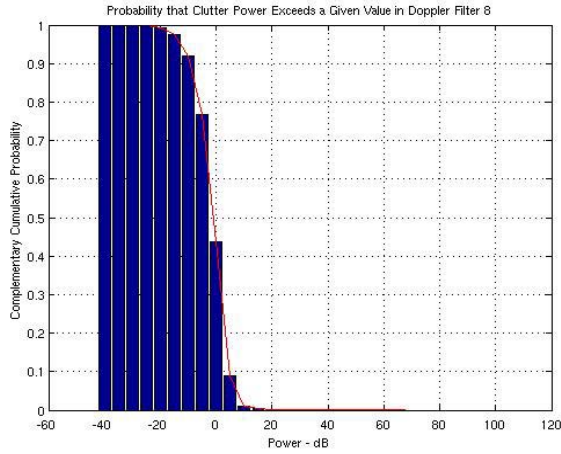


Figure 4. Complementary Cumulative Distribution of Clutter Power at Output of Doppler Filter 8

C. Spatial Correlation of Clutter

A third detection strategy was suggested by observing that the envelope of all the data exhibits some degree of spatial correlation from one range bin to the next, both in the time domain and in the frequency domain. If this spatial correlation is sufficient to predict the clutter magnitude in one range bin based on the clutter in previous range bins, then any large prediction error could be attributed to the presence of an impulsive target.

To estimate the spatial correlation of the clutter data, define the estimated covariance, $\hat{\kappa}_x(m)$, of a sequence of real range samples, $x(n)$, as

$$\hat{\kappa}_x(m) = \frac{1}{N} \sum_{i=0}^{N-1-|m|} (x(i) - \hat{\mu}_x)(x(i+|m|) - \hat{\mu}_x) \quad (1)$$

where N is the length of the sequence, and the estimated mean of $x(n)$ is,

$$\hat{\mu}_x = \frac{1}{N} \sum_{n=0}^{N-1} x(n). \quad (2)$$

The correlation coefficient is defined by,

$$\rho(m) = \frac{\hat{\kappa}_x(m)}{\hat{m}_{20}}, \quad (3)$$

where \hat{m}_{20} is the estimate of the second moment of the data about its mean,

$$\hat{m}_{20} = \frac{1}{N} \sum_{i=0}^{N-1} (x(i) - \hat{\mu}_x)^2. \quad (4)$$

This quantity can be computed in MATLAB using the function call $Var(x, 1)$.

Using all the time domain data, Fig. 5 illustrates the complementary cumulative distribution of the estimated spatial correlation coefficient $\rho(1)$, which measures the correlation between adjacent range bins. The data sequence $x(n)$ in this case is defined to be the magnitude of the entire sequence of 97 complex range samples within one recorded range interval 12 miles long. Fig. 5 shows that all the data exhibits some spatial correlation from range bin to range bin. In fact, one-third of the data has a spatial correlation greater than 0.4.

A possible explanation for the spatial correlation of the data is that the range profile of the regions containing clutter varies slowly. Another possible factor correlating the data is the pulse compression sidelobes of strong clutter extending out in range. In general, the extent of the spatial correlation of the clutter will depend on whether the radar is illuminating land clutter, sea clutter, weather clutter, or a combination of all three. Furthermore, the system parameters of the radar will also affect the measured characteristics of the clutter, including its spatial correlation. For example, the transmitted frequency, the PRF, the transmitted polarization, and the transmitted bandwidth will all have an affect on the clutter statistics. Although the data presented in this paper includes all the clutter detected by the AN/SPS-49 radar at CBD, it does not necessarily represent the clutter that would be detected by a shipboard radar at sea in other parts of the world. Fig. 6 shows that the spatial correlation of the magnitude of the Doppler filter bank output can be even greater than the spatial correlation of the temporal data. For example, almost 40 percent of the output from Doppler filter 1 has a spatial correlation coefficient greater than 0.4, which is a greater fraction of the data than in the time domain.

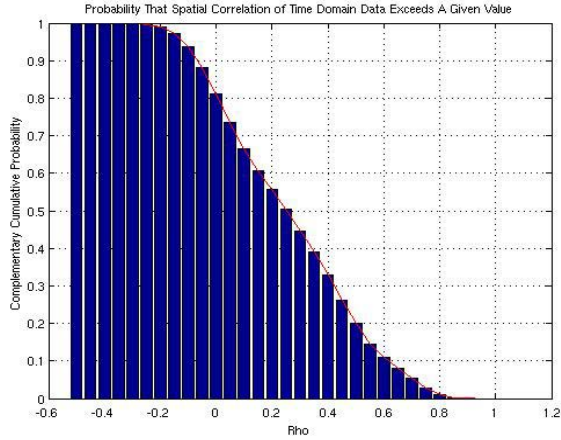


Figure 5. Spatial Correlation of Temporal Data

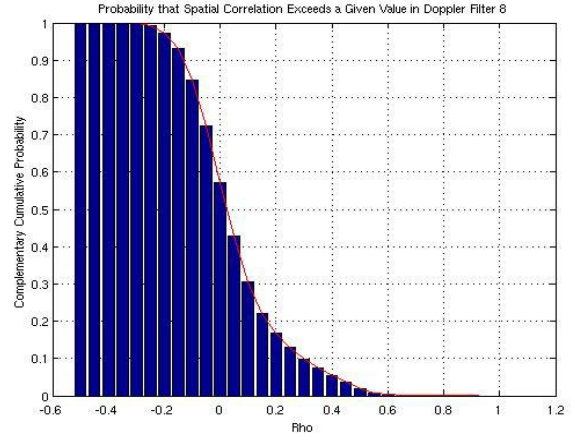


Figure 7. Spatial Correlation at Output of Doppler Filter 8

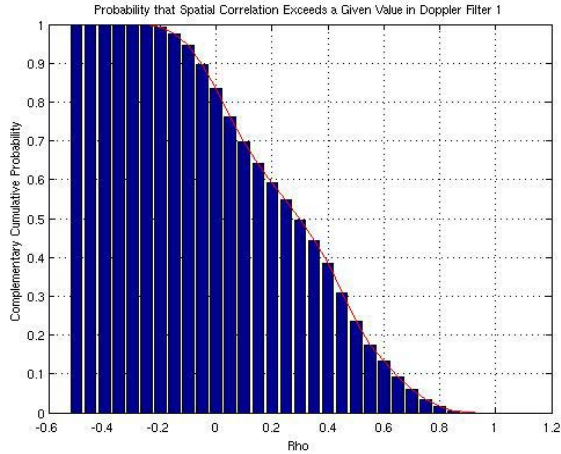


Figure 6. Spatial Correlation at Output of Doppler Filter 1

However, Fig. 7 shows that clutter in the higher Doppler filters, considered to be filters 3 through 14, exhibits very little spatial correlation. For example, only 5 percent of the output from Doppler filter 8 will exhibit a spatial correlation coefficient greater than 0.4. Coupled with the fact that much weaker clutter is present at the output of the higher Doppler filters, suggests that targets can be detected in those filters using conventional techniques, such as Greatest-Of (GO) Cell Averaging (CA) CFAR.

III. PROPOSED ADAPTIVE CLUTTER PROCESSOR

The proposed new adaptive clutter processor (ACP) takes advantage of the spatial correlation exhibited by the clutter envelope in the lower Doppler filters to cancel it. In the higher Doppler filters, targets can be detected without excessive false alarms using a GOCA-CFAR. The basic premise of the ACP is taken from [1] and relies on the use of a LP filter to predict the clutter values in successive range bins. The higher the spatial correlation of the clutter, the better the prediction will be. Any targets present in the range interval will appear at the output of the LP filter as large errors, and can be detected by applying a threshold. Fig. 8 illustrates a block diagram of the proposed receive chain processing.




Figure 8. Proposed Adaptive Clutter Processor

After the LP filter, CFAR block, and the first threshold comparison step, a binary image is formed for each CPI with a one at the range/Doppler cells where threshold crossings occur and zero otherwise. A batch of these binary images is accumulated over several CPIs and integrated using a binary integrator. The output of the binary integrator is then compared to a second threshold. Any range/Doppler cells whose integrated amplitude is greater than zero but less than a second threshold are set equal to one. The result of the second threshold operation is a final binary image of possible tracks created over the course of the binary integration period. Any stationary clutter that leaked through the LP filter should be removed when the second threshold is applied. In the TBD processing block, the final image is processed to remove regions contaminated by noise. A

novel noise removal technique based on morphological filters is applied to yield a clean image void of noise with the true target tracks exposed, as well as possibly a few false tracks. Morphological filters perform nonlinear operations on images using set theoretic concepts and can be particularly effective in removing noise.

A. Overview of Lattice Predictor Filters

This section provides a little background on linear prediction. For a more thorough treatment consult [2]. Let X_{n-1} denote the M -dimensional space spanned by the samples, $x(n-1), x(n-2), \dots, x(n-M)$. Then the predicted signal value $\hat{x}(n|X_{n-1})$ is defined by,

$$\hat{x}(n|X_{n-1}) = \sum_{k=1}^M w_{f,k}^* x(n-k), \quad (5)$$

where $w_{f,k}$ are the forward prediction weights used to combine the past samples and * denotes conjugation. In the present context, the desired output response, $d(n)$, of the prediction filter equals $x(n)$. The forward prediction error, $f_M(n)$, equals the difference between the input sample $x(n)$ and its predicted value $\hat{x}(n|X_{n-1})$,

$$f_M(n) = x(n) - \hat{x}(n|X_{n-1}) = x(n) - \sum_{k=1}^M w_{f,k}^* x(n-k). \quad (6)$$

The subscript M signifies the order of the filter, which is the number of unit delays necessary to store the samples used to make the prediction. A filter that operates on the set of samples $x(n-1), x(n-2), \dots, x(n-M)$ to produce the forward prediction error $f_M(n)$ at its output is called a forward prediction-error filter.

We may also predict backwards in time by using the subset of M samples $x(n), x(n-1), \dots, x(n-M+1)$ to make a prediction of the sample $x(n-M)$. Let X_n denote the M -dimensional space spanned by $x(n), x(n-1), \dots, x(n-M+1)$. We make a linear prediction of the sample $x(n-M)$ as given by,

$$\hat{x}(n-M|X_n) = \sum_{k=1}^M w_{b,k}^* x(n-k+1), \quad (7)$$

where $w_{b,k}$ are the backward prediction weights. The backward prediction error $b_M(n)$ equals the difference between the desired response $x(n-M)$ and the linear prediction of it. Therefore,

$$b_M(n) = x(n-M) - \sum_{k=1}^M w_{b,k}^* x(n-k+1). \quad (8)$$

A filter that operates on the set of samples $x(n), x(n-1), \dots, x(n-M+1)$ to produce the backward prediction error $b_M(n)$ at its output is called a backward prediction-error filter.

Consider the lattice prediction filter of order M with M stages shown in Fig. 9. The output of the top channel is the forward prediction error and the output of the bottom channel is the backward prediction error. If the input to the lattice predictor is the output of a Doppler filter, then the ACP is the top channel of the lattice predictor, and its output is ideally a clutter-free signal. At the beginning and trailing edges of a region of contiguous clutter, or where there is a target present, there will be a large error at the output of the ACP. However, in regions where the clutter magnitude is correlated, one would expect to see small errors at the output of the ACP, primarily determined by the amount of thermal noise present.




Figure 9. Lattice Predictor Filter

The reflection coefficients, k_m , used in the lattice predictor are computed using,

$$k_m(n) = \frac{-2N_{m-1}}{D_{m-1}}, \quad m = 1, 2, \dots, M \quad (9)$$

where

$$N_{m-1}(n) = \sum_{i=1}^n \lambda^{n-i} f_{m-1}(i) b_{m-1}^*(i-1), \quad (10)$$

$$D_{m-1}(n) = \sum_{i=1}^n \lambda^{n-i} \left[|f_{m-1}(i)|^2 + |b_{m-1}(i-1)|^2 \right].$$

The weighting constant λ is a scalar between 0 and 1 used to de-emphasize distant data and emphasize recent data. Reducing λ has the effect of forgetting past data at a quicker rate, which allows the filter to adapt faster in a non-stationary environment. For this study, λ was set equal to 0.9, after using a process of trial and error to find the best value.

Fig. 10 illustrates the output of the LP filter for a CPI taken at random from a random scan. Only clutter and thermal noise are present in this CPI. The blue curve represents the input to the LP filter, which is the Doppler filter output with clutter present. The red curve is the output of the LP filter, which is the forward prediction error. The plot shows that the magnitude of the clutter at the output of the LP filter has been substantially reduced.




Figure 10. Output of LP Filter Showing Clutter Cancellation

Fig. 11 shows the signal before and after the LP filter in a random interval when there is a 10 dB target present in the data. Because the target is an isolated spike, the LP filter does not predict its existence and the target appears in range bin 80 at the output of the LP filter as an error. The target can be detected with much fewer false alarms after the LP filter because the clutter in the range interval has been significantly canceled. Furthermore, there is minimal signal loss on the target due to the LP filter processing.

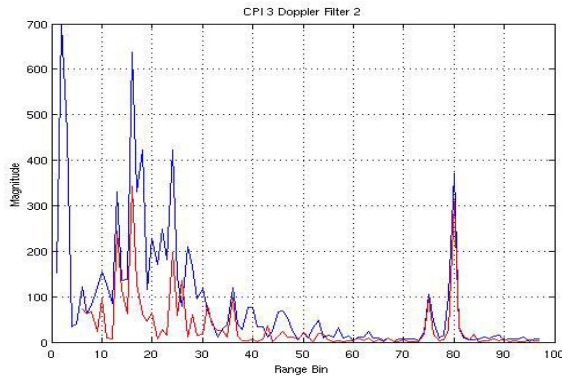


Figure 11. Output of LP Filter With Target Present

B. CFAR

The output of Doppler filters 3 through 14 was processed using conventional techniques because clutter was not as pervasive as in the lower Doppler filters. Specifically, the output of those Doppler filters was passed through a GOCA-CFAR circuit and then compared to a detection threshold. Fig. 12 illustrates the operation of the GOCA-CFAR circuit. A measurement window on either side of the Cell Under Test (CUT), designated by a T in the figure, averages 4 samples of the background noise level. One guard cell on either side of the CUT, designated by G , isolates any strong target sidelobes from the measurement windows. The greatest-of the noise estimates, designated by N , is used to normalize the amplitude of the CUT. Then the normalized signal samples are compared to a detection threshold. Any range bins with amplitudes greater than the detection threshold are set equal to one, and any range bins with amplitude less, are set equal to zero.

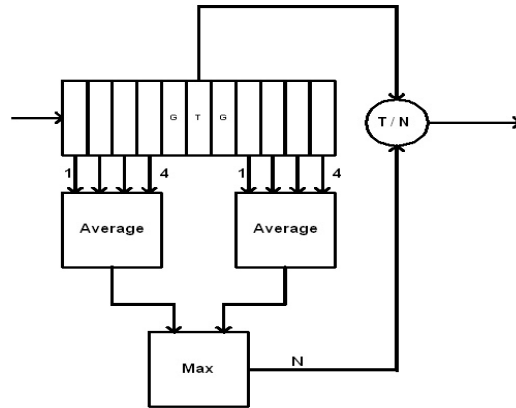


Figure 12. Operation of GOCA-CFAR Circuit

One problem with GOCA-CFAR processing is that if a large clutter spike is present in one of the measurement windows, it will greatly skew higher the estimated background noise level, especially if the measurement windows are short, and consequently suppress any target return in the CUT. A more robust CFAR design for certain clutter environments is the Order Statistic (OS) CFAR shown in Fig. 13 [3]. The OS-CFAR operates by ranking the signal samples in the measurement windows based on their amplitudes. It then selects the k th largest sample as an estimate of the background noise level. The utility of this approach is that by avoiding consideration of the largest amplitudes in the measurement windows, the estimated noise level is less likely to be skewed if spiky outliers are present in the data.

At the output of the OS-CFAR, a decision is made as to whether a target is present by comparing the signal magnitude, or power, to a detection threshold. The threshold is set by multiplying the estimated interference level at a range bin by a predetermined scalar, t . If the signal magnitude, or power, exceeds the threshold, then a target may be present at the threshold crossing. Setting the detection threshold and computing the probability of false

alarm for the OS-CFAR can be done accurately if the probability density of the input samples is known. Otherwise, Monte Carlo methods may be used to select the threshold empirically, which was done for this study.

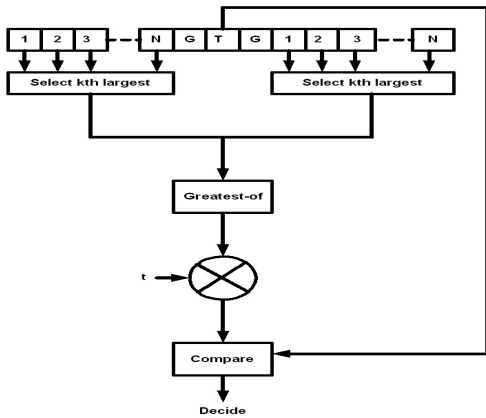


Figure 13. Operation of OS-CFAR Circuit

C. Morphological Filtering

The track-before-detect processor is essentially a morphological filtering operation on the binary image of range/Doppler threshold crossings produced at the output of the binary integration and second threshold comparison step. A morphological filter uses set theoretic operations, such as unions and intersections, to filter sizes and shapes from an image as opposed to frequencies. In this case, the objective is to filter straight lines embedded in the range/Doppler image corresponding to target tracks. More information on morphological filtering can be found in [4]. A very useful property of morphological filters for radar applications is that they are extremely effective in removing pepper noise and were used in the TBD processing block to denoise the final binary image and expose the target tracks. An acquisition tracker could also have been used to reduce or eliminate the noise in the image.

IV. RESULTS

A. Ideal Output

To test the performance of the ACP, 3 targets were injected into the radar traveling at 20, 70, and 120 knots. To establish the ideal performance of a receiver for comparison to the ACP, the targets were initially injected at a high Signal-to-Noise Ratio (SNR) of 20 dB while the radar transmitted into a dummy load instead of radiating into space. Fig. 14 illustrates the binary image created in this instance. The tracks of the 3 targets are clearly visible.

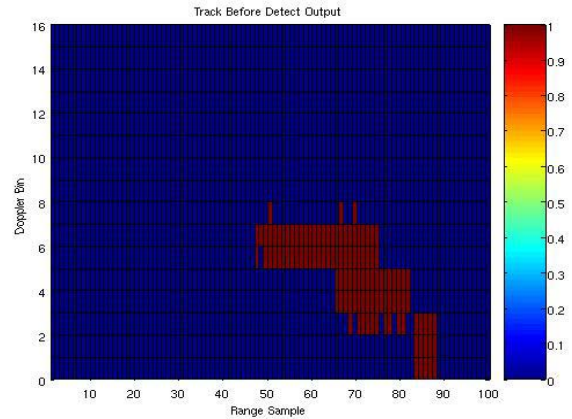


Figure 14. Output of Ideal Receiver

B. Output of Conventional Processor

For purposes of further gauging the performance of the ACP, the more conventional processor depicted in Fig. 15 was used to process the data. This processor is essentially a Doppler filter bank with Dolph-Chebyshev weighting, followed by a two-pass GOCA-CFAR. The two-pass GOCA-CFAR is very similar to the one-pass GOCA-CFAR shown in Fig. 12, with the exception that large signal amplitudes are excised from the data before estimating the background noise and clutter level. The threshold for the two-pass GOCA-CFAR was set to a level corresponding to a probability of false alarm equal to $1e-6$.

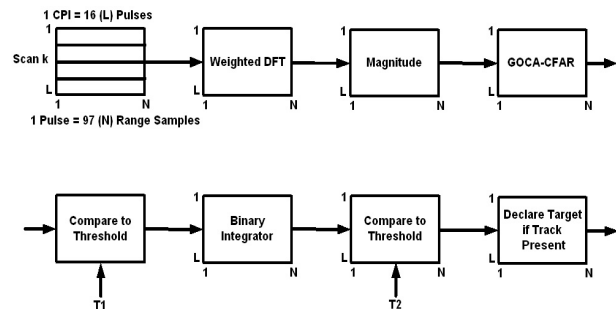


Figure 15. Conventional Signal Processor

Fig. 16 shows the output of the conventional signal processor when the 3 targets following the same trajectories shown in Fig. 14 were injected into the radar with a SNR of 10 dB. The plot shows that none of the targets, but plenty of clutter, was detected. This illustration underscores the difficulty of detecting targets in a spiky clutter environment using typical GOCA-CFAR schemes.

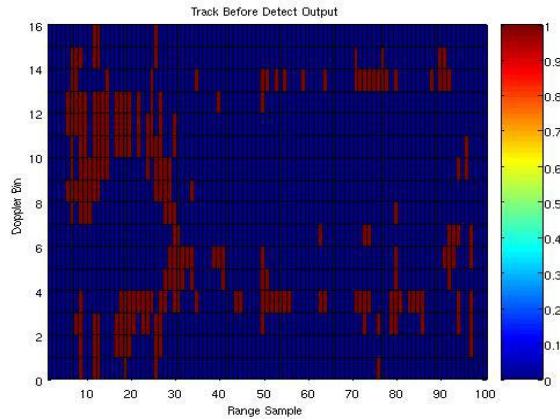


Figure 16. Output of Conventional Signal Processor

C. Output of ACP

To test the ACP, the 3 targets flying the same trajectories as before were injected at random beam positions and ranges, with a SNR equal to 10 dB into the radar. Fig. 17 shows the binary range/Doppler image of threshold crossings after the binary integration of CPIs, but before the TBD processing block.

The image in Fig. 17 has a lot of pepper noise which consists of spurious outputs that span 1, 2, or up to a few range bins. This type of noise can be effectively removed by morphological filtering operations, the results of which are shown in Fig. 18. Notice that part, if not all, of the true target tracks appear in the ACP output, as opposed to the output of the conventional processor shown in Fig. 16, which shows none of the target tracks.

A few false tracks appear in Fig. 18 that are artifacts of the morphological filtering operations and of incomplete clutter cancellation. Nevertheless, the number of false tracks is low and could be reduced by observing the targets over a longer period of time or by using excision logic to eliminate the tracks that are too long to be created in a particular Doppler bin during the observation time. Fig. 19 illustrates one more sample of the ACP output selected at random.

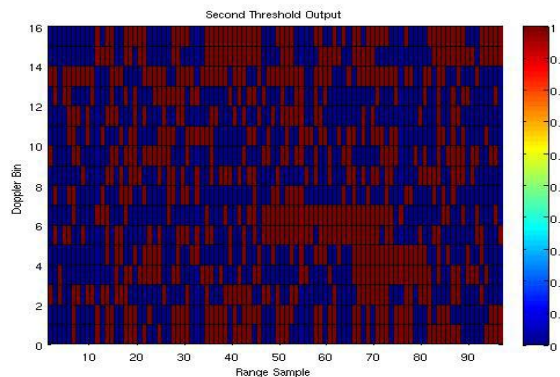


Figure 17. Before Morphological Filtering




Figure 18. Output of ACP After Morphological Filtering

The results presented in this section do not definitively characterize the performance of the ACP in a general clutter environment with many different target scenarios. However, this preliminary investigation suggests that linear prediction and morphological filtering are two processing techniques that have the potential to improve detection performance in a clutter environment and they merit further investigation.

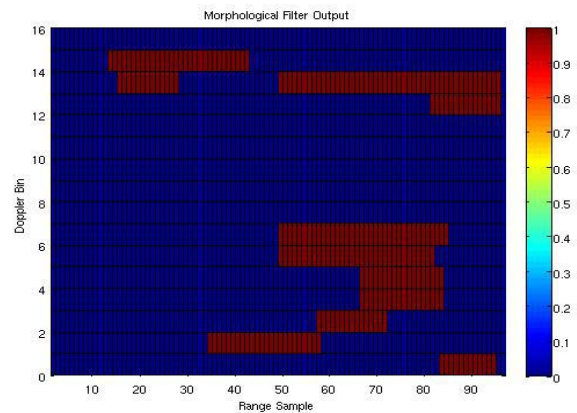


Figure 19. Output of ACP – Example 2

REFERENCES

- [1] Xu, Feng, Hao, "Adaptive Radar Clutter Suppression", Department of Electrical and Computer Engineering, California Polytechnic University, Technical Report MTS 0-933957-28-9.
- [2] Haykin, Simon, Adaptive Filter Theory, Third Edition, Prentice Hall, 1996.
- [3] Rifkin, R, "Analysis of CFAR Performance in Weibull Clutter", IEEE Transactions on Aerospace and Electronic Systems, Vol. 30, No. 2, April 1994.
- [4] Giardina, C. R. and Dougherty, E. R., Morphological Methods in Image and Signal Processing, Prentice Hall, 1988.